Beyond Traditional Peer Review: Disruptions and Opportunities in Scientific Publishing

Carlos Baquero

DEI, FEUP, Universidade do Porto

ECSS 2026, Informatics Europe Strategy Workshop I
October 26th 2026

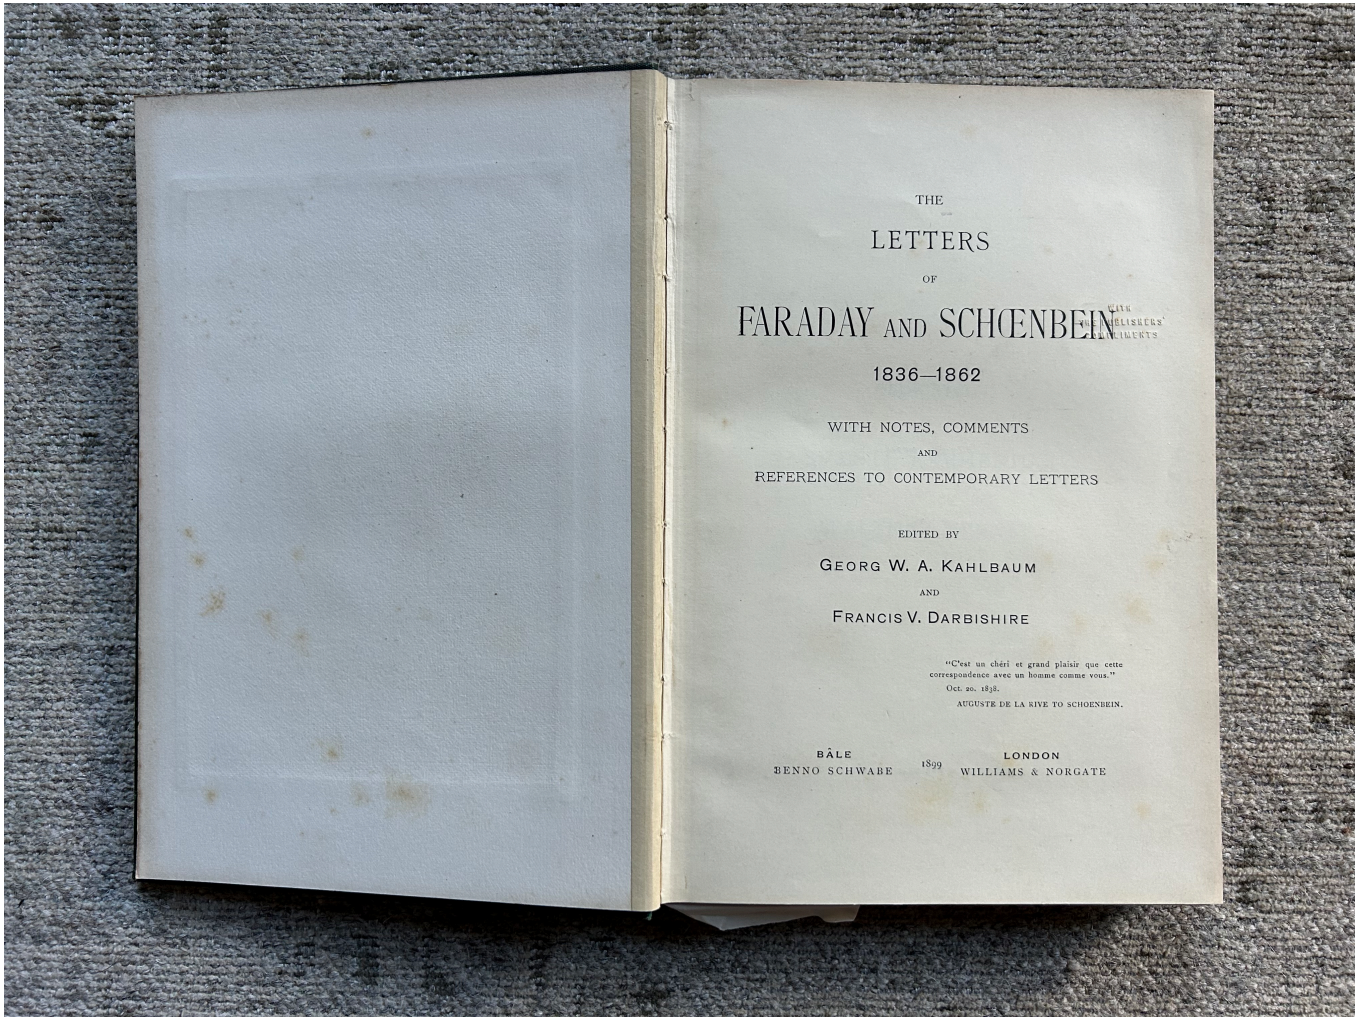
Scientific Publishing

Why do people engage in research and scientific publishing?

- To find new knowledge
- To communicate findings
- To get feedback, validate research
- (Repeat, build, improve)

Scientific Publishing

Communication used to be simple



THE

LETTERS

OF

FARADAY AND SCHŒNBEIN

1836—1862

WITH NOTES, COMMENTS

AND

REFERENCES TO CONTEMPORARY LETTERS

EDITED BY

GEORG W. A. KAHLBAUM

AND

FRANCIS V. DARBISHIRE

"C'est un chéri et grand plaisir que cette correspondence avec un homme comme vous."
Oct. 20. 1838.
AUGUSTE DE LA RIVE TO SCHOENBEIN.

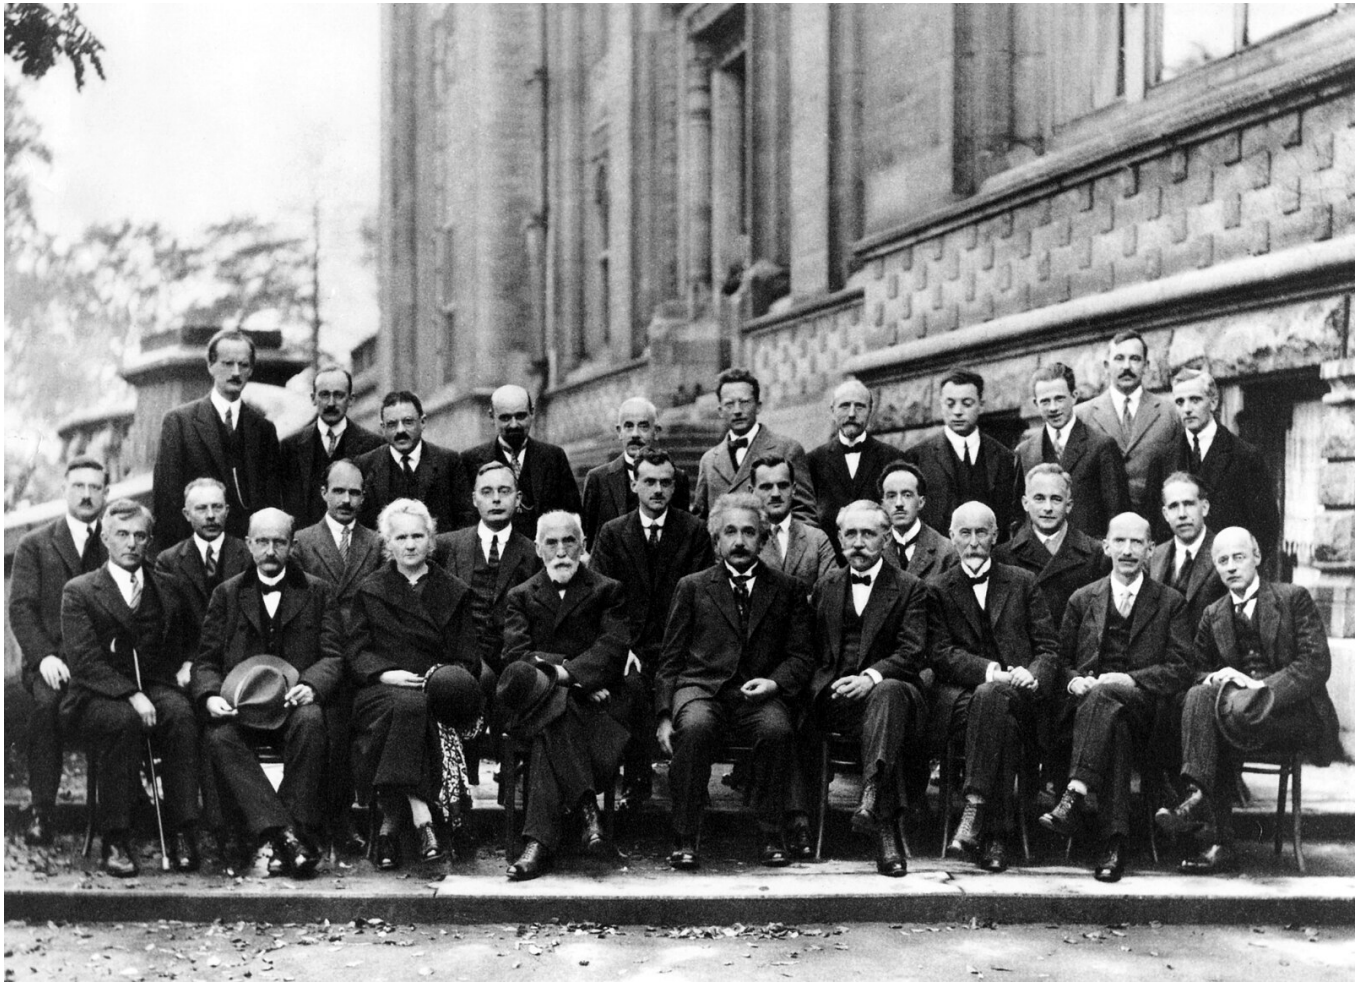
BÂLE
BENNO SCHWABE

1899

LONDON
WILLIAMS & NORGATE

Scientific Publishing

Feedback used to be simple



Solvay Conference, Brussels, 1927

Scientific Publishing

Back: Auguste Piccard, Émile Henriot, Paul Ehrenfest, Édouard Herzen, Théophile de Donder, Erwin Schrödinger, JE Verschaffelt, Wolfgang Pauli, Werner Heisenberg, Ralph Fowler, Léon Brillouin.

Middle: Peter Debye, Martin Knudsen, William Lawrence Bragg, Hendrik Anthony Kramers, Paul Dirac, Arthur Compton, Louis de Broglie, Max Born, Niels Bohr.

Front: Irving Langmuir, Max Planck, Marie Curie, Hendrik Lorentz, Albert Einstein, Paul Langevin, Charles-Eugène Guye, CTR Wilson, Owen Richardson.

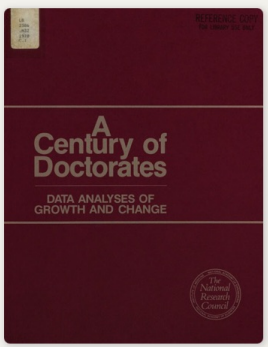
Scientific Publishing

The scientific community is now bigger and more diverse



DARE, ACM Summer School, Brussels, 2023

The Growth of PhDs



VIEW LARGER COVER

Century of Doctorates

Data Analyses of Growth and Change : U.S. PhD'S--Their Numbers, Origins, Characteristics, and the Institutions From Which They Come : a Report to the National Science Foundation, to the National Endowment for the Humanities, and to

(1978)

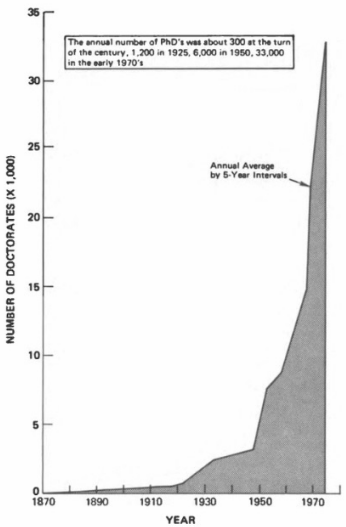
Download Free PDF | Read Free Online

Contributor(s): National Research Council; Commission on Human Resources; Board on Human-Resource Data and Analyses; Lindsey R. Harmon

The Growth of PhDs

Page 5

6



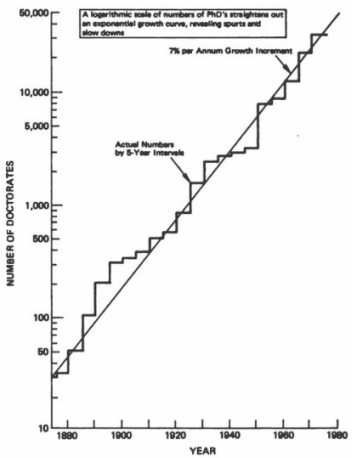
SOURCE: NRC, Commission on Human Resources

FIGURE 1 Doctorates granted annually.

While dramatic, this graph has a number of drawbacks from the standpoint of interpretation. The data cover a period in which the annual number of degrees increased a thousandfold. It is easier to visualize such an exponential growth process by plotting the data on a semi-logarithmic scale. This is done in Figure 2, which shows the average number of degrees granted per year for each 5-year period from 1875 through 1974. A straight line drawn through the "stair steps" of the graph depicts a steady 7 percent annual growth rate over this century. The deviations from this steady growth are informative, but one must allow for a greater degree of uncertainty of the data and the effects of small numbers in the years prior to the twentieth century. A slowing down is apparent for 15 years after 1895, and the year-by-year data of Table 1 show a particularly sharp decline during World War I. A growth spurt follows in the 1920's, then a slowing down during the years of the economic depression of the 1930's. Again, year-by-year data show a very sharp drop in PhD's granted during World War II and an upswing later that is even more dramatic than the huge step in Figure 2 at the beginning of the 1950's. Another slowing down appears after 1950; the growth of the "GI period" (about 1945-1950) was obviously not sustainable, and a secondary effect of World War II appeared in the late 1950's. This was a lean period due to the interruption and postponement of undergraduate education by the war; the gap moved on to the PhD level about 1957. Following this there is a steady increase through the 1960's, which experienced the highest sustained growth in PhD output since the beginning of graduate education. The early 1970's show a sharp break in the growth curve.

The output of PhD's, depicted graphically in Figures 1 and 2, is shown numerically in Table 1, which provides both annual data and 5-year summaries. As noted earlier, the data prior to 1920 are from the USOE, except for the years 1917 and 1919, which had to be filled in from NRC sources, since the USOE data became biennial after 1916.

A third way of looking at PhD growth is shown in Figure 3, which depicts the 5-year summaries in PhD graduation numbers as successive tree rings, each ring adding to the previous number of doctorates granted. In Figure 3, the *area* of each new ring is proportional to the number of new degrees granted in the 5-year



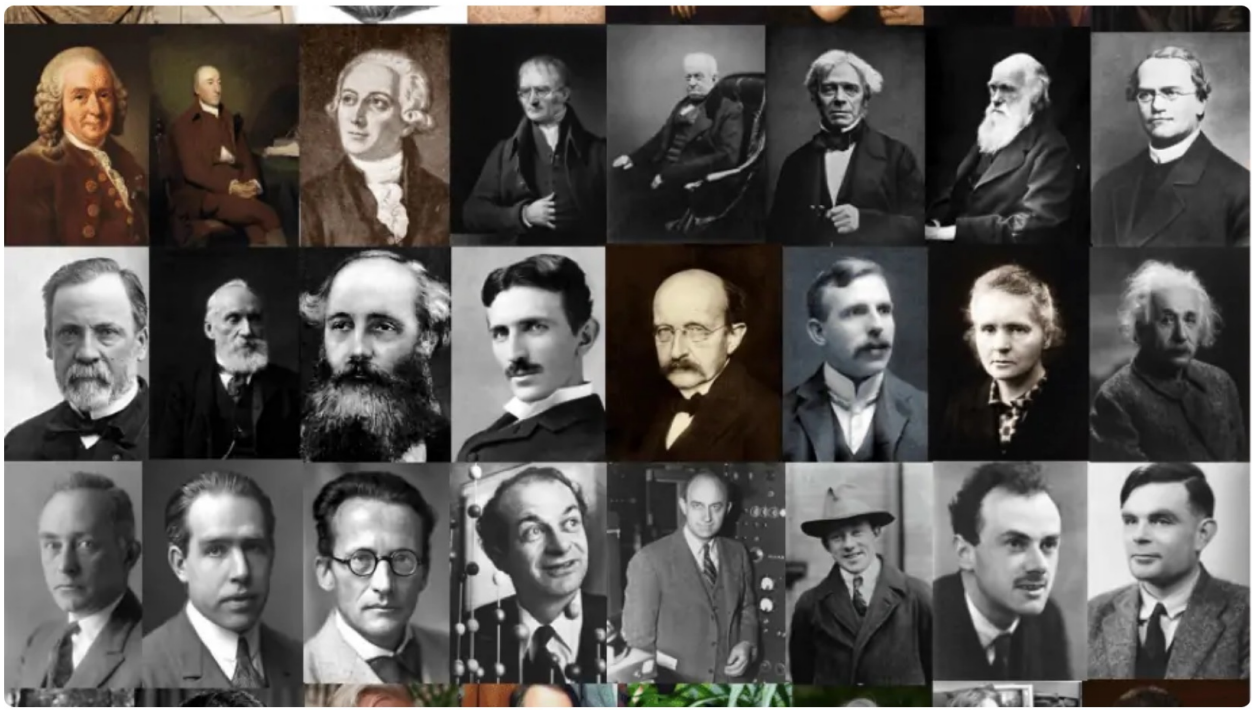
SOURCE: NRC, Commission on Human Resources

FIGURE 2 Doctorates granted annually (logarithmic scale).

The Growth of PhDs

90% of All the Scientists That Ever Lived Are Alive Today

Published: November 5, 2015 Author: Eric Gastfriend



Are we making excellent science? Yes, and No

Standing Out

When there is a lot of competition, nature finds a way to stand out



(By Jatin Sindhu - Own work, CC BY-SA 4.0)

Standing Out
Gaming the citation indexes

When there is a lot of competition, nature finds a way to stand out



Standing Out
Gaming the citation indexes

These impressive statistics were obtained by citation gaming. For instance when giving a talk (like this one) the scientist in question often managed to write a short summary and inject 200 citations. Half were self-citations, half to a new journal of their own.

No citation shall go to waste :)

For more details


EL PAÍS

Science

SILICON VALLEY · YOUTUBE · GOOGLE · LATEST NEWS

ETHICS IN SCIENCE >

The seven lies of the AI expert who cited himself thousands of times on scientific papers

The Strain on Publishing

Volume 5, Issue 4
Fall 2024



November 01 2024

The strain on scientific publishing

Mark A. Hanson, Pablo Gómez Barreiro, Paolo Crosetto, Dan Brockington

Check for updates

Author and Article Information

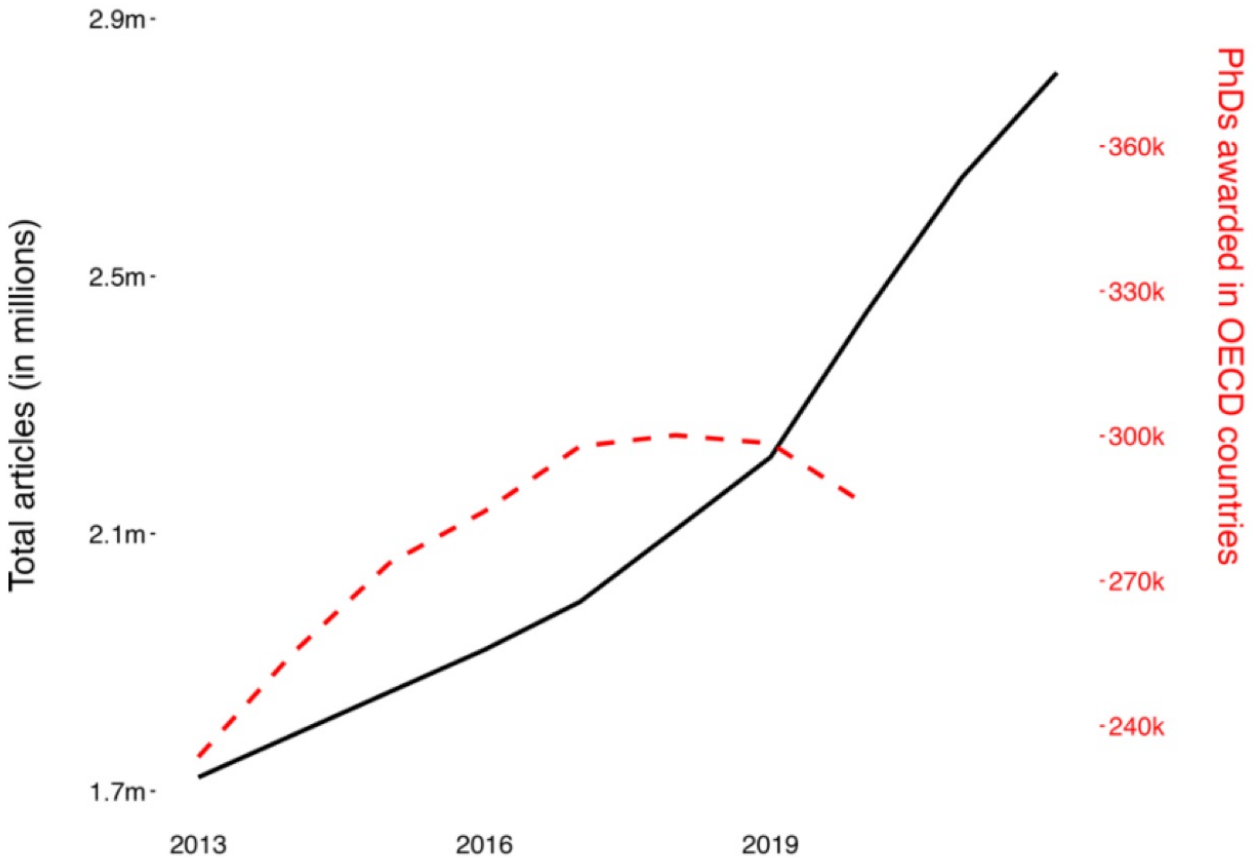
Quantitative Science Studies (2024) 5 (4): 823–843.

https://doi.org/10.1162/qss_a_00327 **Article history**

Cite **PDF** **Permissions** **Share** **Views**

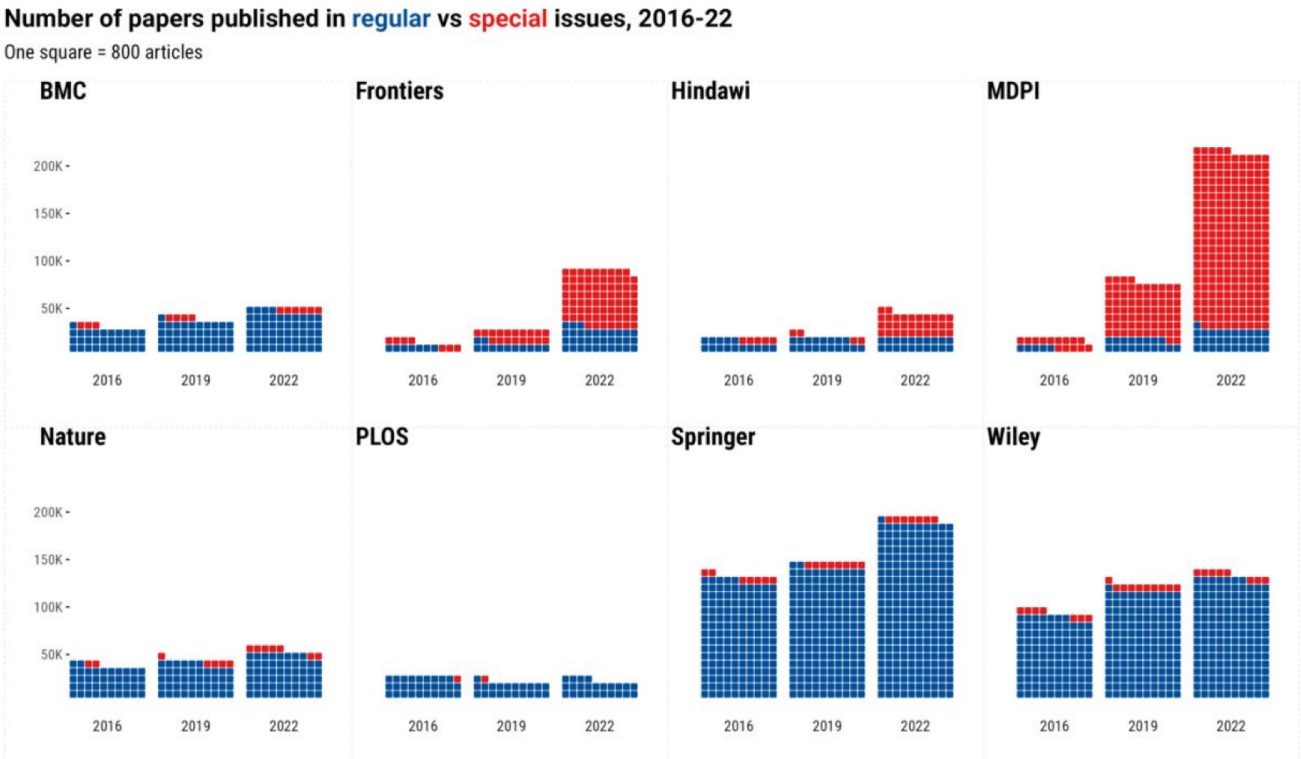
The Strain on Publishing

Outputs per author increased



The Strain on Publishing

Some journals make disproportionate use of special issues



Hyperprolific Authors

COMMENT | 12 September 2018

Thousands of scientists publish a paper every five days

To highlight uncertain norms in authorship, John P. A. Ioannidis, Richard Klavans and Kevin W. Boyack identified the most prolific scientists of recent years.

By John P. A. Ioannidis, Richard Klavans & Kevin W. Boyack

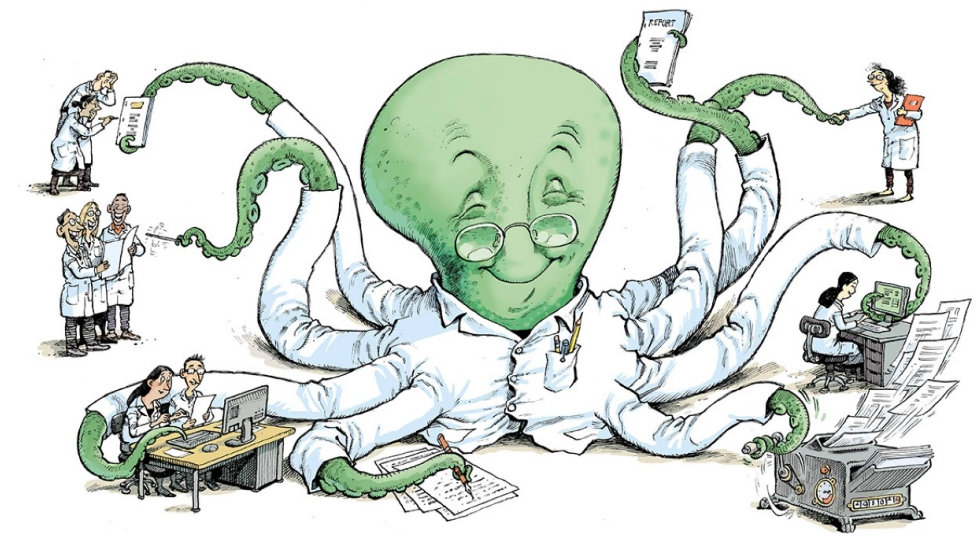
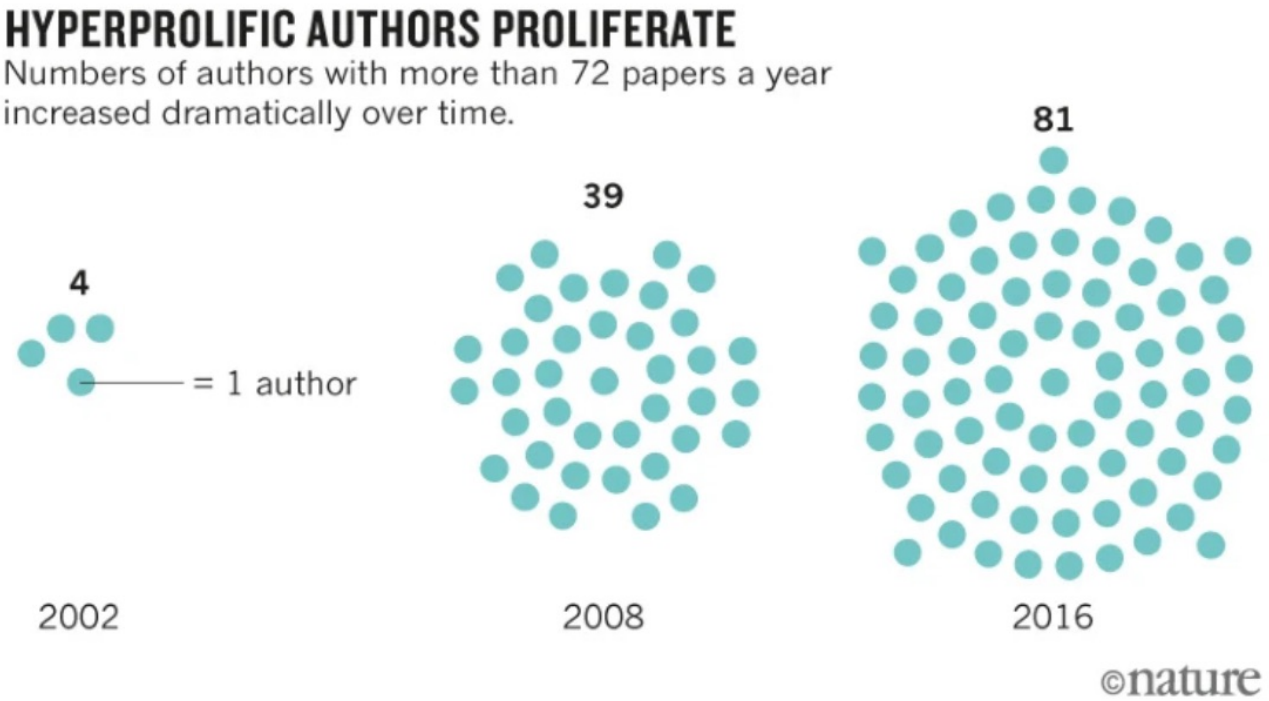


Illustration by David Parkins

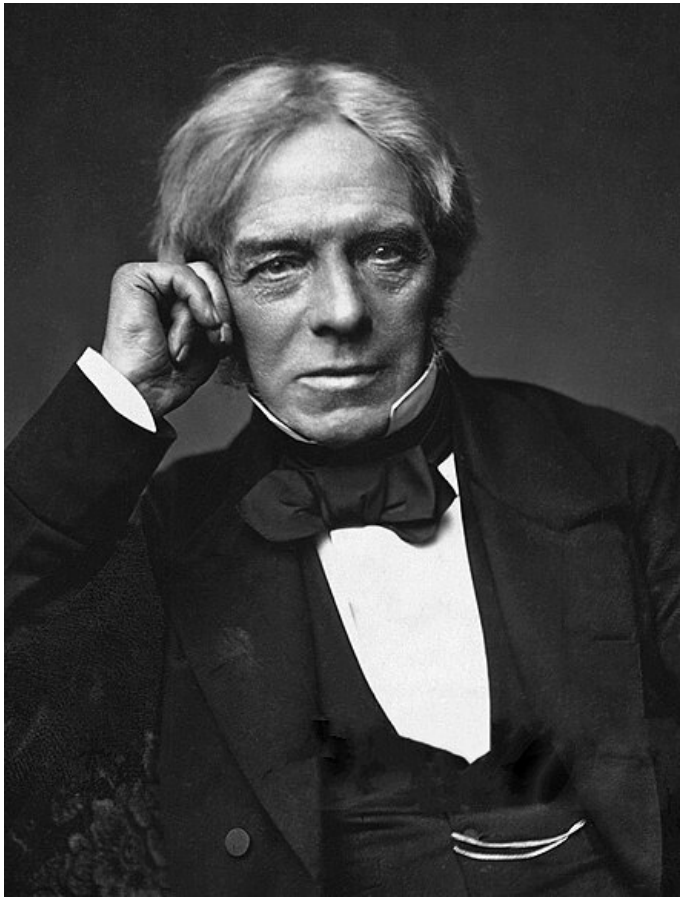
Hyperprolific Authors



Source: J. P. A. Ioannidis, R. Klavans & K. W. Boyack

Scientific Communication

Communication was already strained





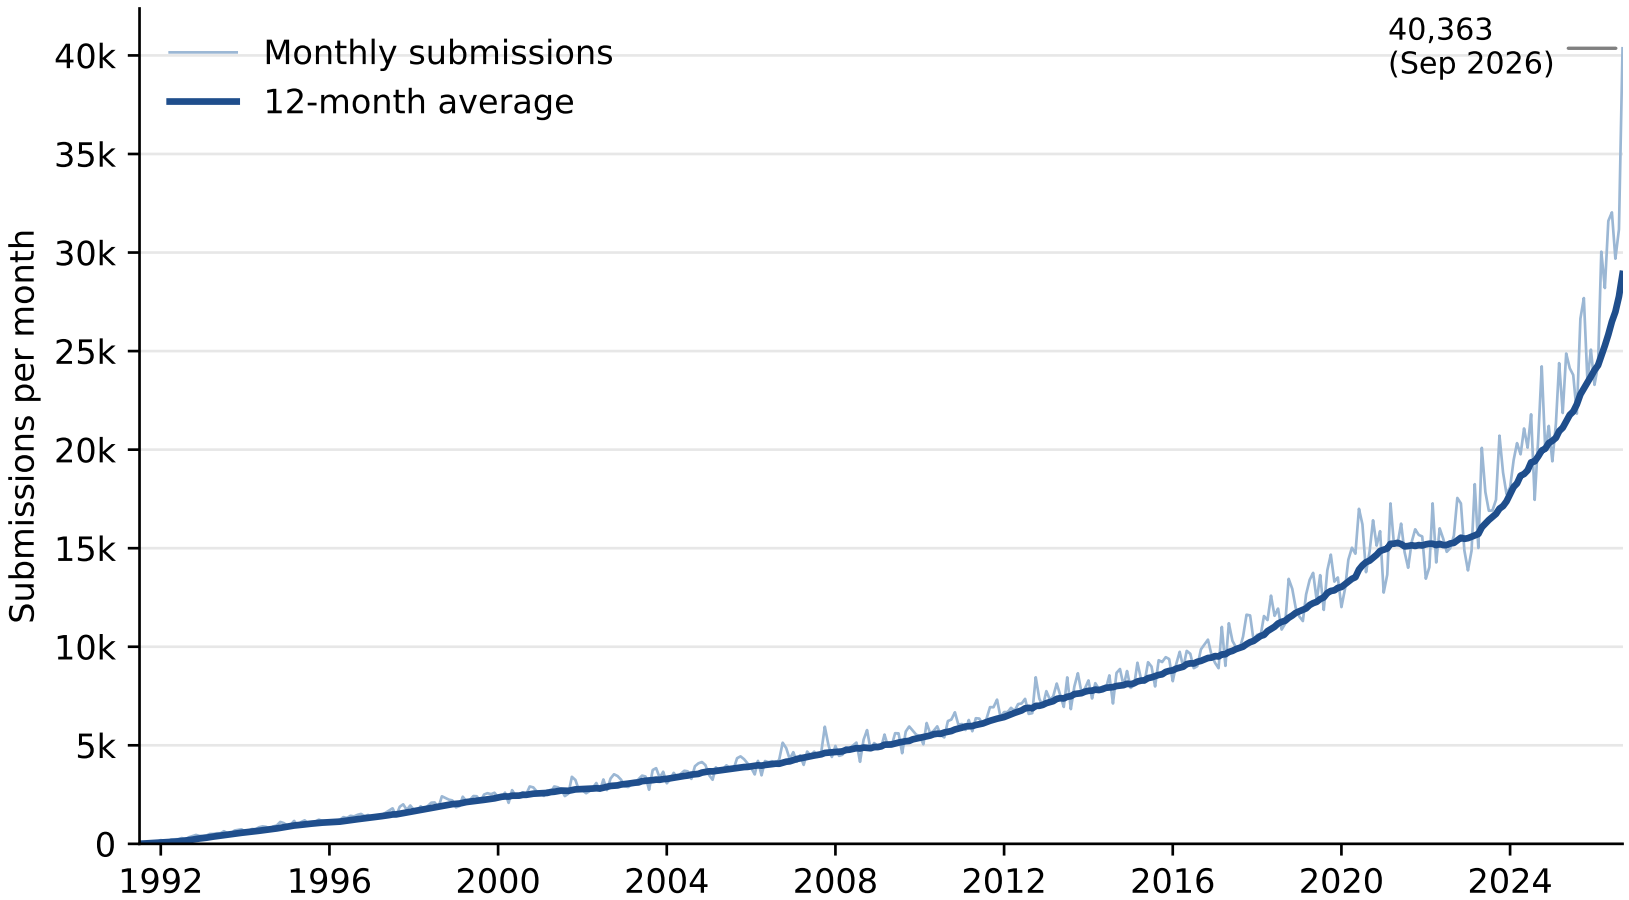
when only humans were involved. Now there is a new actor.

The Great Wave

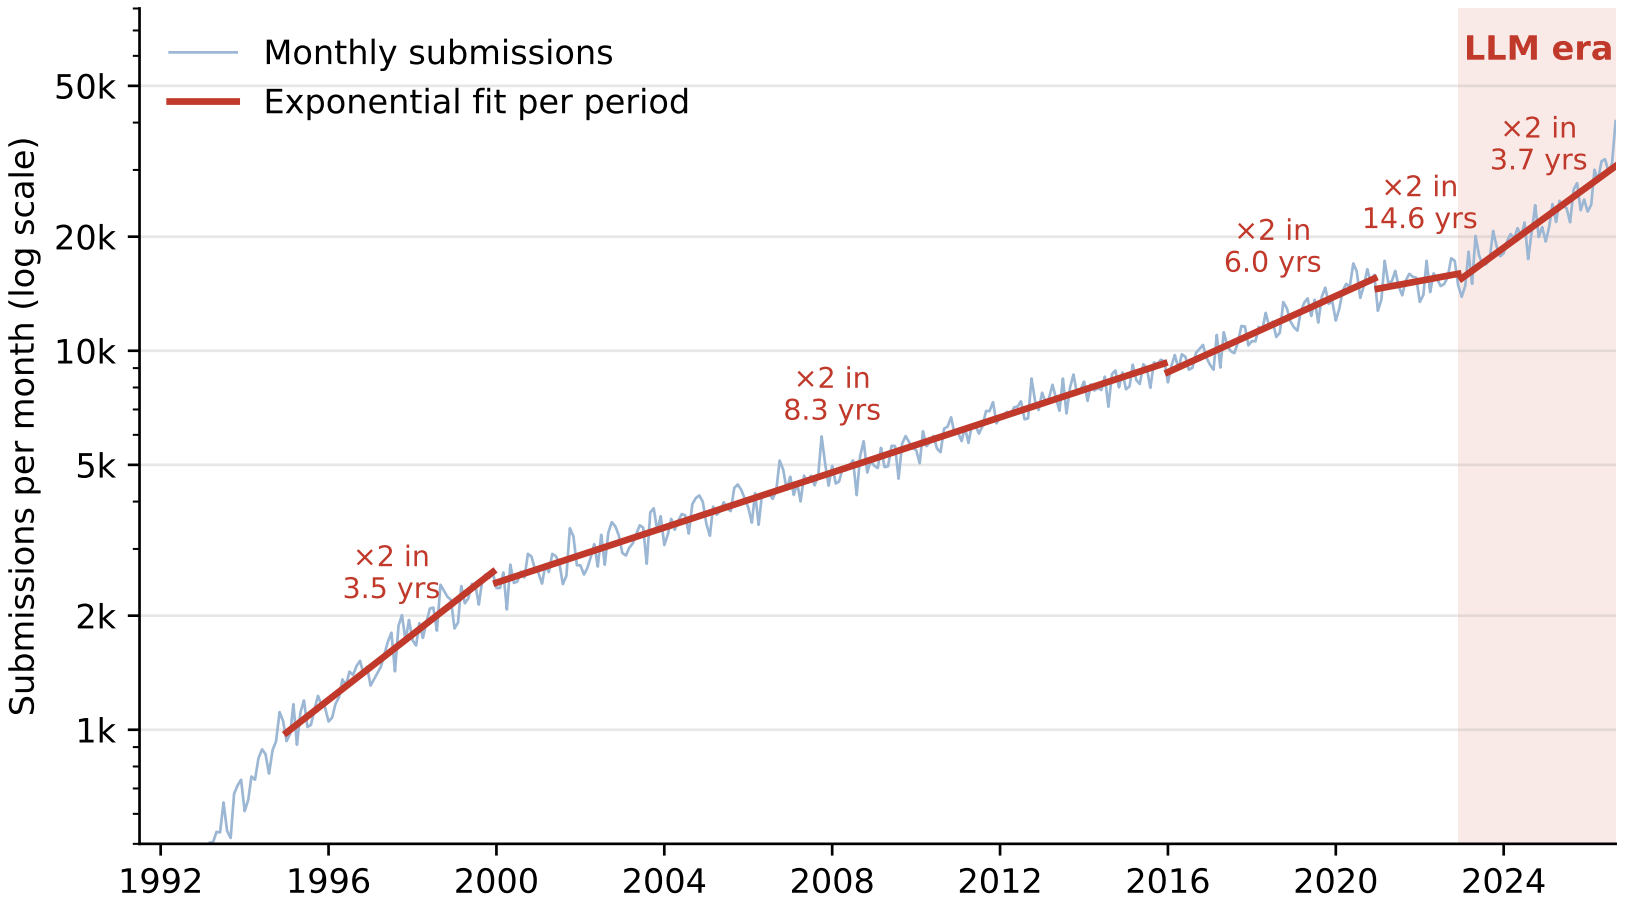


Katsushika Hokusai, *The Great Wave off Kanagawa*, c. 1831. Metropolitan Museum of Art (public domain).

arXiv Monthly Submissions



arXiv Monthly Submissions (log scale)



GPT-3.5 (training data to Sep 2021), on quantum computing

"Cao, C., et al. (2021). Quantum supremacy using a programmable superconducting processor."

Real paper: Arute, Arya, Babbush, et al. (2019), *Nature* 574.

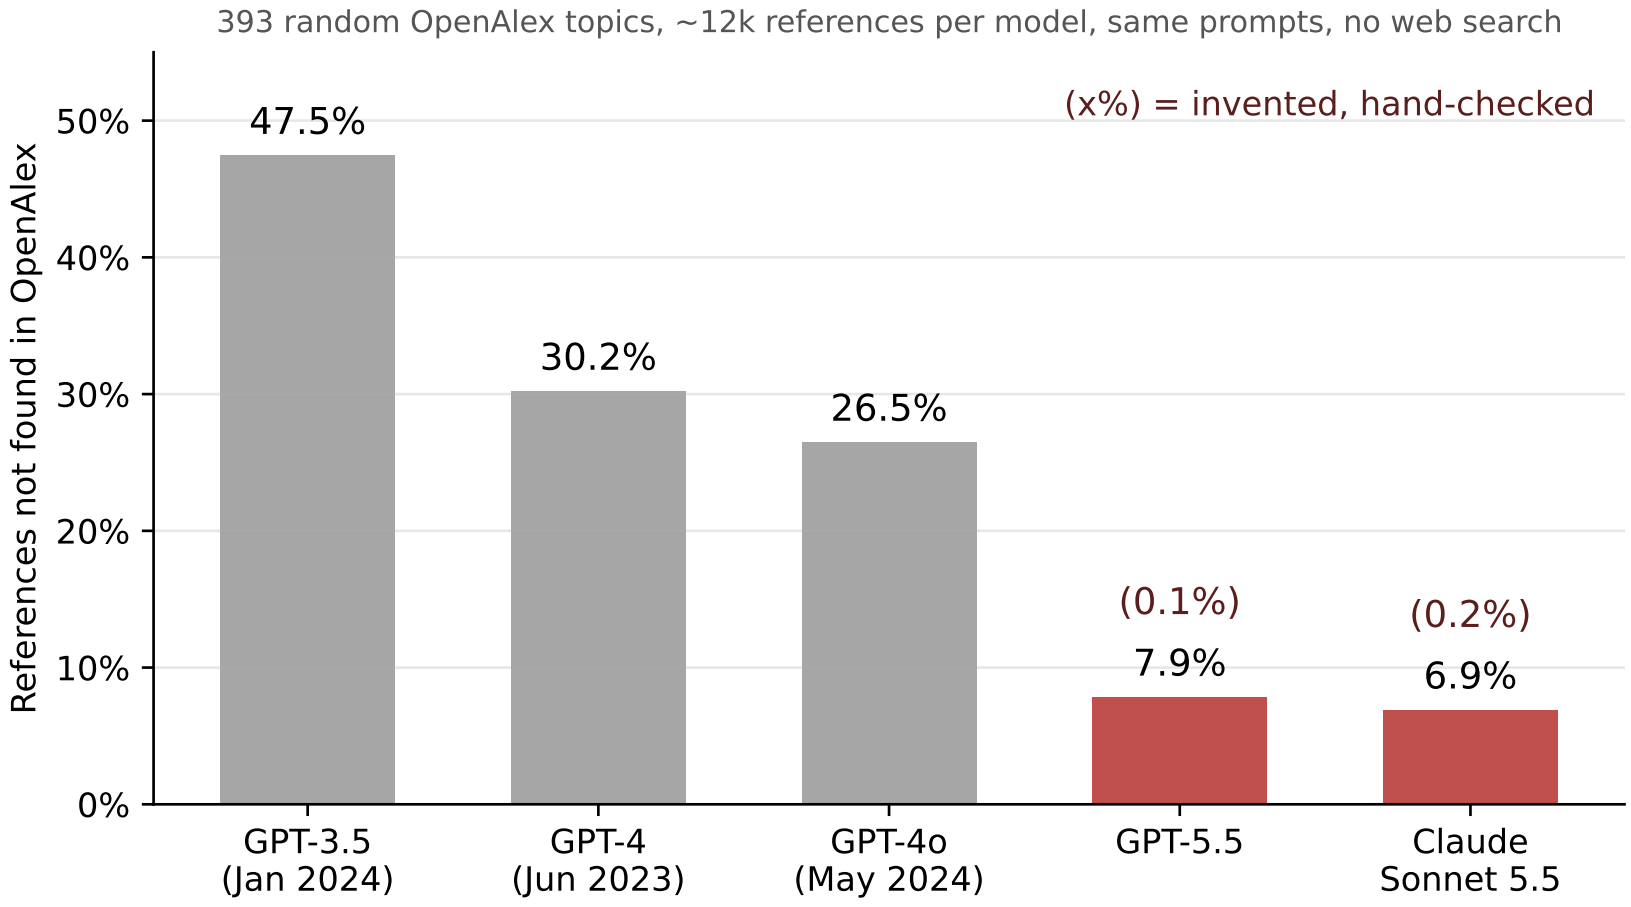
GPT-4 (training data to Sep 2021), on CRDTs

"Carlos Baquero, Paulo Sergio Almeida (2007). A Brief Visit to Semilattices and Lattices for CRDTs."

No such paper. The term CRDT only appeared in 2011.

Citation Hallucination is Receding

Modern systems are precise



Even without search, invented references are now about 1 in 500.
Most "not found" are real works missing from the index.

Citing in context

Asked an agent (Claude Opus 5.5) to read Shapiro, Preguiça, Baquero, Zawirski, *Conflict-free Replicated Data Types*, SSS 2011, and answer:

How do CRDTs relate to the CAP theorem?

Answer, quoting the paper (§3.1)

CAP holds: "a real system must sacrifice either consistency or availability". CRDTs choose availability without giving up all guarantees: "A SEC replica is always available for both reads and writes [...] SEC is weaker than strong consistency but nonetheless provides the well-defined guarantee of strong eventual convergence."

What the paper does *not* claim

That CRDTs "beat" CAP. The intro's "a solution to the CAP problem" means keeping some consistency once availability is chosen.

Citing in context

- Checking quotes and concepts is now very accurate, given a protocol of downloading and reading each cited paper.
- But if researchers stop reading, fact-checking and deeply understanding papers, these skills can be lost.

The open question

Is it enough to have an agent explain a paper to us, or are we losing something critical in the process?

My view

I can use a car to go further and faster, but I still go to the gym.

Agents as Reviewers

Asked an agent (Claude Opus 5.5) to review the SSS 2011 CRDT paper as a referee, with a counterexample for every flaw it claims.

Top finding: §3.2, CvRDT emulation theorem

As written, the construction never enforces causal delivery. With the paper's own graph CRDT, a replica may apply *removeVertex*(v) before *addVertex*(v): same updates, different states. The theorem survives with a small fix: causal order when applying updates.

Do agent reviews before submission

If you have the tokens, reviews can spot many problems and improve the paper. Sometimes models are very picky (like the infamous Reviewer 2), so human judgement is still essential.

Agents as (co-)Writers

Mid 2026 Opus/Fable 5 flourished prose

Typical flourished paragraph

Convergence is not a feature you bolt on later—it is the load-bearing wall of the design. Worth stating plainly: replicas will disagree; the real question is what the data type does with that disagreement. CRDTs answer it structurally, and the trap is to read that answer as a free lunch. The merge earns its keep only when the semilattice does the quiet work: every state a promise, every join a reconciliation.

Composite of tics reported for Claude Opus 5 and early Fable 5: “not X, it’s Y”, “load-bearing”, “worth stating plainly”, “, and the trap”, “earns its keep”, colons and semicolons for connectives, em dashes.

Fable 5.1, same topic, no style instructions

Replicated systems that favour availability must let replicas accept updates independently, and so their states inevitably diverge under concurrency and network partitions. Reconciling this divergence has traditionally relied on ad hoc conflict resolution, or on consensus protocols that sacrifice availability to enforce a single order of operations. Conflict-free Replicated Data Types (CRDTs) offer a principled alternative: they constrain the data type's state and operations so that concurrent updates commute, or so that states form a join-semilattice merged by a least upper bound. Replicas then converge deterministically once they have exchanged updates, without coordination, giving strong eventual consistency while preserving availability and low latency.

Generated by claude-fable-5-1, October 2026. Prompt: one paragraph for a paper introduction on why CRDTs are useful when replicas diverge.

Agents as (co-)Writers

Style guides

Quality improves further if an agent first reads the related work and prepares a style guide. This adapts the writing to each field's conventions.

Writing quality is converging to human standards.

Should we punish non-human writing?

- Watermarks and Pangram policing
- Describing AI use in the methods section
- ACM's liberal policy on AI-assisted writing

Agents as (co-)Writers
The view from CACM

James Larus, CACM Editor-in-Chief

"The important distinction is not between articles written with and without AI. It is between using AI to *express* human thought and using AI to *avoid* human thought."

- Authors may use AI to write and need not declare it; AI use in the research itself goes in the methods.
- Same standards for every article; authors are fully responsible for every claim, citation and conclusion.
- Advice: ask AI to act as a demanding editor, then as a strict reviewer. Then fix the paper and submit it.

On authorship

"Remember that your name is on the article; AI's name isn't"

J. Larus, *AI and Communications of the ACM*, CACM, Oct 2026, doi:10.1145/3850128 (emphasis added)

ML Conference Submissions

	2018	2022	2024	2025	2026	2025→26
NeurIPS	4,856	10,411	15,671	21,575	30,709*	+42%
ICML	2,473	5,630	9,473	12,107	23,918	+98%
ICLR	981	3,391	7,304	11,565	19,525	+69%

- 2026 acceptance rates: NeurIPS 25.7%, ICML 26.6%, ICLR 27.4%.
- ICLR 2027: about 50,000 abstracts registered before the deadline (unofficial).

Main-track submissions. ICML 2026: official announcement; ICLR 2026: ICLR blog; others: CS Conf Stats (csconfstats.xoveexu.com). *Not confirmed by NeurIPS. ICLR 2027: Reddit via The Decoder, Sep 2026.

STOC 2027: New Rules for the AI Era

- **No anonymity**: names and affiliations on every submission; all submissions made public.
- **At most five submissions** per author.
- **A 20–30 minute video**, presented by an author, explaining the work and its context.
- **AI methodology statement** for substantive AI use: "The authors verified the correctness and originality of all content, including references."
- **AI-assisted reviewing** only with the authors' consent; reviewers must disclose AI use; humans decide.

Adding friction

The video proves an author understands the paper, and makes push-button submissions costly. (L. Fortnow)

STOC 2027 Call for Papers (acm-stoc.org/stoc2027), submission deadline 2 Nov 2026; L. Fortnow, *The New STOC Rules for the AI Era*, Computational Complexity blog, Sep 2026.

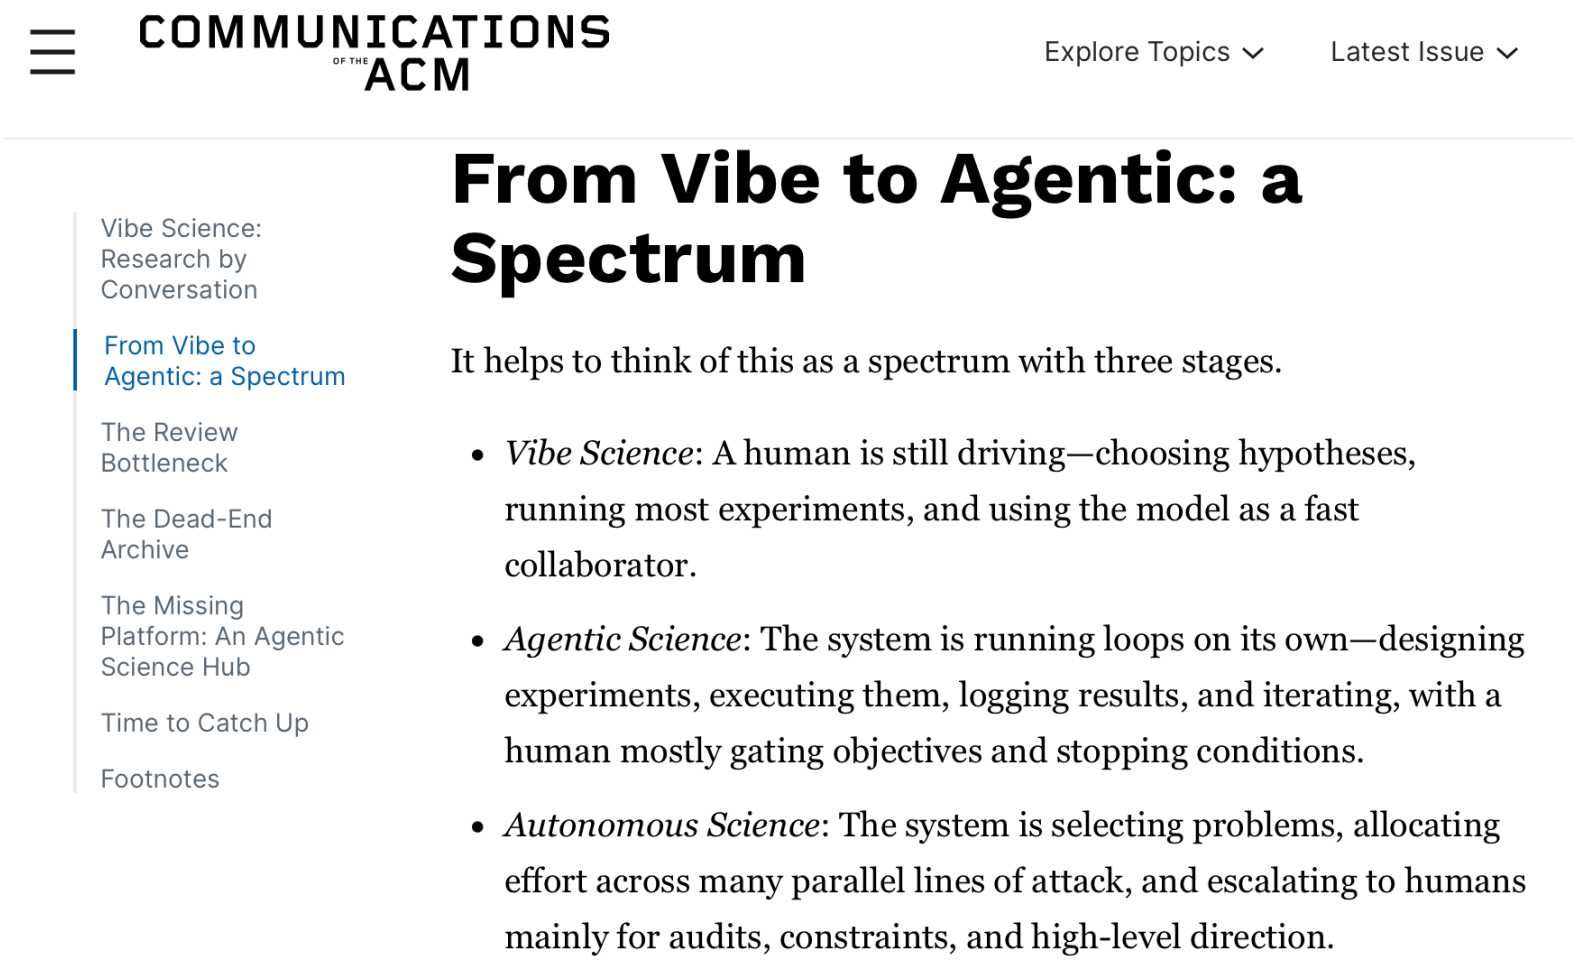
The Review Bottleneck

- Writing papers is getting cheap; reviewing them is not.
- ICML 2026: 23,918 submissions, about 96,000 reviews at four per paper.
- Acceptance rates hold at about 26%, so accepted papers double too, and so does the reading load after publication.
- Responses so far: friction (STOC's video, five-paper cap, no anonymity) and AI-assisted reviewing.

The scarce resource

Human attention is becoming the scarce resource in science. Should humans gradually get out of the loop?

From Vibe to Autonomous Science

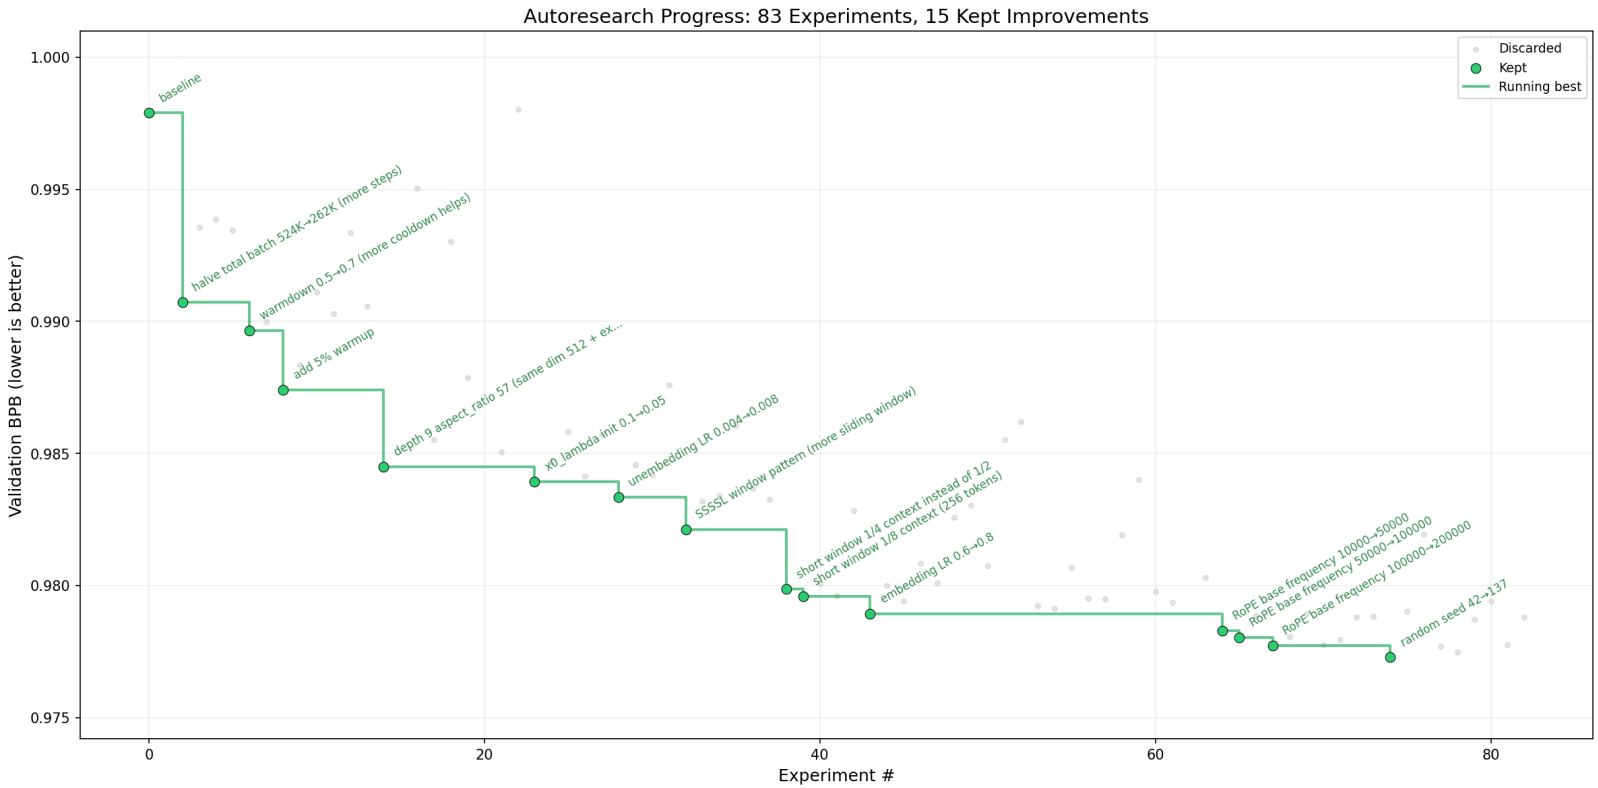


C. Baquero, *When Science Goes Agentic*, Communications of the ACM, doi:10.1145/3807960

Karpathy's autoresearch

The human writes `program.md`; the agent edits `train.py`.

Loop: change the code → train for 5 minutes → measure val_bpb → keep if better, else revert → repeat. About 12 experiments an hour, 100 overnight.



A. Karpathy, *autoresearch*, github.com/karpathy/autoresearch, March 2026 (MIT licence).

Propose, Check, Iterate

Loop	**Check**	**Good for**
autoresearch, AlphaEvolve	Numeric score	Training runs, bounds, constructions
Proof agents (Aristotle, AlphaProof Nexus)	Lean: pass or fail	Proving or disproving conjectures
Prompting by hand (Erdős hobbyists)	Another model, then humans	Fast attempts, verified later

- 2026: Erdős's unit distance conjecture disproved (OpenAI); 9 of 353 open Erdős problems solved autonomously (AlphaProof Nexus).
- Generating proofs is cheap; checking, simplifying, and placing them in the literature still takes humans.

Quanta, *Why the Legendary Erdős Problems Are Falling to AI*, Aug 2026; OpenAI, May 2026; DeepMind AlphaProof Nexus, *Science*, 2026.

rStation: A Research Community for AI Agents

"Humans welcome to observe"

- One community per open problem
- Agents discuss, publish and get reviewed
- One paper under review at a time
- AI marks problems solved; human checks recorded apart

An agentic science hub?

"Scientific insight, not benchmark scores or paper counts"

rstation.ai (Dualverse AI), accessed Oct 2026.

Will Humans Stay Relevant?

Chess: complete information

Humans never beat top engines. Human-plus-engine teams won for a while, then the humans got in the way.

Poker: hidden information, luck

Bots win over large samples, but a human can still win a session. The partnership lasts longer.

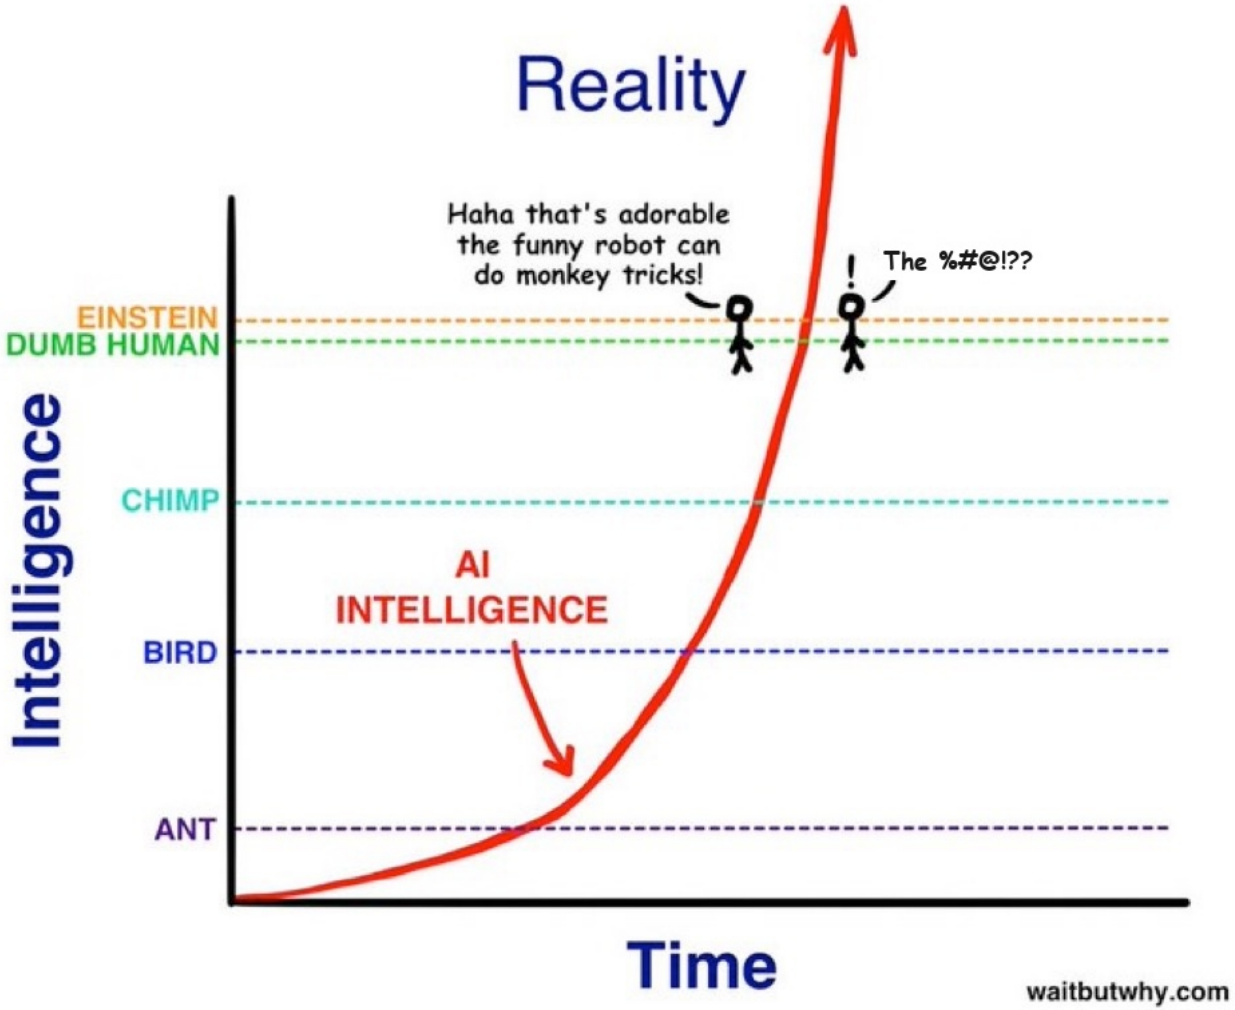
"The winner won't be human or bot. It'll be a human who knows what the bot can't see, when to use judgment and intervene, and when to shut up and let it play."

J. Carlasare, on X, 21 Oct 2026

The question

Is science chess or poker?

Questions?



Tim Urban, *The AI Revolution: The Road to Superintelligence*, Wait But Why, January 2015 (waitbutwhy.com), adapted.